



University of Pittsburgh

Learning Research and Development Center

AD-A218 827

Difficulty in Learning to Read Speech Spectrograms: The Role of Visual Segmentation

Gareth Gabrys

Learning Research and Development Center
University of Pittsburgh

Technical Report No. LRDC/PITT/IMP-1

Cognitive Science Program

Office of Naval Research

Contract No. N00014-86-K-0361

The research reported in this document was supported by the Cognitive Science Program, Office of Naval Research, under the contract number noted above. The views stated or implied by this report are not necessarily those of the United States Government or the United States Navy, which has neither reviewed nor endorsed the reported findings.

Reproduction in whole or part is permitted for any purpose of the U.S. Government.

Approved for public release; distribution unlimited.

Difficulty in Learning to Read Speech Spectrograms: The Role of Visual Segmentation

Gareth Gabrys
Learning Research and Development Center
University of Pittsburgh

Technical Report No. LRDC/PITT/IMP-1
Cognitive Science Program
Office of Naval Research
Contract No. N00014-86-K-0361

The research reported in this document was supported by the Cognitive Science Program, Office of Naval Research, under the contract number noted above. The views stated or implied by this report are not necessarily those of the United States Government or the United States Navy, which has neither reviewed nor endorsed the reported findings.

Reproduction in whole or part is permitted for any purpose of the U.S. Government.

Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a REPORT SECURITY CLASSIFICATION			1b RESTRICTIVE MARKINGS		
2a SECURITY CLASSIFICATION AUTHORITY			3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
2b DECLASSIFICATION/DOWNGRADING SCHEDULE					
4 PERFORMING ORGANIZATION REPORT NUMBER(S) UPITT/LRDC/ONR/IMP-1			5 MONITORING ORGANIZATION REPORT NUMBER(S)		
6a NAME OF PERFORMING ORGANIZATION Learning Research & Development Center, Univ. of Pittsburgh		6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Cognitive Science Program Office of Naval Research (Code 1142CS) 800 North Quincy Street		
6c ADDRESS (City, State, and ZIP Code) Pittsburgh, Pennsylvania 15260			7b ADDRESS (City, State, and ZIP Code) Arlington, VA 22217-5000		
8a NAME OF FUNDING/SPONSORING ORGANIZATION		8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0361		
8c ADDRESS (City, State, and ZIP Code)			10 SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO 61153N	PROJECT NO RR04206	TASK NO RR04206-01
11 TITLE (Include Security Classification) Difficulty in Learning to Read Speech Spectrograms: The Role of Visual Segmentation					
12 PERSONAL AUTHOR(S) Gareth Gabrys					
13a TYPE OF REPORT Technical		13b TIME COVERED FROM _____ TO _____		14 DATE OF REPORT (Year, Month, Day) 1990, February, 7	
15 PAGE COUNT 30					
16 SUPPLEMENTARY NOTATION					
17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD 05	GROUP 09	SUB-GROUP	Training difficulty; context-dependent; visual segmentation; spectrograms; speech spectrogram reading; learning difficulties; task demands; salience.		
19 ABSTRACT (Continue on reverse if necessary and identify by block number) This work examines possible sources of training difficulty encountered by learners of speech spectrogram reading. Such difficulty has been attributed to the context-dependent nature of the visual segmentation of spectrogram patterns (Liberman et al, 1968), and suggestions by researchers of other difficult skills (Biederman & Shiffrar, 1983) have also implicated visual segmentation. In both cases, the discriminations necessary to distinguish important parts can be easily made once identified, but are enormously difficult to discover. The experiments presented here used a pseudo-spectrogram reading task which varied the segmentation rules subjects were required to discover. Experiment 1 found that considerable learning difficulty could be produced by this task, but confounded the source of that difficulty among several factors. The second experiment attempted to identify the sources of the difficulty. Segmentation was found to contribute significantly. The salience of the important cues, and, potentially, the demands of the learning task were also found to increase the difficulty of discovering important visual distinctions. These results are discussed.					
20 DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21 ABSTRACT SECURITY CLASSIFICATION unclassified		
22a NAME OF RESPONSIBLE INDIVIDUAL Dr. Susan Chipman			22b TELEPHONE (Include Area Code) 202/696-4318		22c OFFICE SYMBOL ONR 1142CS

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

Block NO. 19 - Abstract- Continued:

with respect to the skill of spectrogram reading and theories of perceptual attention learning.

A-1

A-1

Difficulty in Learning to Read Speech Spectrograms: The Role of Visual Segmentation

When acquiring a new perceptual skill, a learner is usually faced with the problem of learning to recognize new features and discovering which combinations of features form meaningful patterns. In X-ray reading, for example, a student must learn which features indicate normal tissue and which indicate diseased tissue. Such learning is cognitive: the visual system breaks up the visual array into parts and recognition responses occur to learned features, but cognitive processing and training are required to make decisions about which parts are important and how they combine to form higher-level patterns.

Theories of perceptual learning have characterized this cognitive processing as an hypothesis-and-test procedure (Levine, 1975; Trabasso & Bower, 1968) which results in the building of pattern detectors (Kahneman, 1973; Chase & Simon, 1973). More recently, attention has focused on the types of preferences or heuristics which may be required to constrain hypothesis search in complex displays (Michalski, 1983; Medin, Wattenmaker & Michalski, 1987). One type of constraint the cognitive system must make is where to draw object boundaries, i.e., which parts belong together as objects. Characteristics such as spatial relations, overlap, proximity, and shading differences may play a role in determining object coherence (Triesman, 1986). For certain perceptual skills, however, such segmentation decisions can create difficulties. For example, in x-ray pictures, brightness corresponds to the density of tissue rather than any reflective property (Squire, 1988). Hence, visual contours and separations may not correspond to organ or tissue boundaries. For example, if two organs of equal density abut, no contour will appear between them. A radiology student needs to learn a new way of segmenting an x-ray picture to identify the locations of organs and other tissue groups.

The present research is concerned with learning difficulties which may result when visual segmentation does not correspond to object segmentation. Its focus is on a skill which until recently was considered extremely difficult if not impossible to learn: speech

spectrogram reading. Much of the difficulty in spectrogram reading has been attributed to problems in segmenting the display. The goal of this research was, first, to show that learning difficulty could be produced by violating segmentation assumptions, and, second, to look at how segmentation interacts with other stimulus and task variables.

A speech spectrogram is a graph of the energy in different frequency components of speech over a short sampled time. Its two axes represent frequency and time, and the darkness of a small region represents the amount of sound energy at the frequency and time matching its coordinates. When real-time spectrographic displays were first developed, it was hoped that people, especially the hearing impaired, could be taught to recognize speech by seeing it. However, learning to identify speech from this graphical display has proven to be difficult, requiring both an understanding of acoustic-phonetics and many hours of practice. Potter, Kopp, and Green (1947), in one of the earliest efforts toward such training, taught a group of subjects to identify important acoustic features in spectrograms and then had them try to communicate with each other using a real-time spectrographic display. They found that the time to learn the most common words spoken by a single person increased linearly with practice, at the rate of about 4 words per hour. That is, prior learning did not aid the learning of new words. A similar learning rate was found by Greene, Pisoni, and Carrell (1984), who had naive subjects learn to identify spectrograms of 50 words made by a single speaker. The subjects began with four words and were gradually given additional sets of four words over 22 sessions. After about 13 sessions the subjects were able to learn the new items with few errors and show a fair amount of transfer to a new list of words by the same speaker (91.3%) and the original word list spoken by a different speaker (76%). These studies have been viewed optimistically as demonstrating that people can be trained to recognize visual speech. However, the studies are limited by their use of speech from a single speaker, or by their focus on learning of individual words which would not generalize well to continuous speech.

More impressive has been the effort of Dr. Victor Zue, who has taught himself to read spectrograms of continuous speech.

independent of speaker, with a high level of accuracy (Cole, Rudnick, Zue, & Reddy, 1980). Zue systematically studied spectrogram patterns for one hour per day over several years. This extensive practice along with his expertise in acoustic-phonetics has discovered features and rules which enable him to identify phoneme segments with an accuracy of about 85%. The features Zue uses are spectral patterns unique to individual phonemes, but he augments simple detection of these features with knowledge of coarticulation effects, which can distort the features, and a knowledge of phonotactic constraints in the English language. Zue has also been successful in identifying the rules he uses to recognize phonemes and in teaching others to use the rules to read spectrograms with much less practice (40 hrs vs 2000 hrs) (Cole & Zue, 1980).

But what is the original source of the difficulty which limited subjects in early studies to small vocabularies, and required 2000+ hours of training plus acoustic-phonetic knowledge on the part of Victor Zue? In an article entitled "Why are speech spectrograms hard to read?", Liberman et al (1968) identify the major reason for this learning difficulty as the context-dependent nature of the acoustic signal. How a sound is articulated, and hence how it appears on a spectrogram, depends on what other sounds are made immediately before and after it. A vowel following a /d/ will look different from one following a /g/. Context dependency leads to a special learning difficulty because of the inherent difference between the way the visual and auditory systems segment the acoustic pattern. To the visual system, a vowel followed by a stop consonant appears as a wide dark band beside a narrow dark band with a blank space in between, i.e., two distinct objects. However, the auditory segmentation of those two sounds is more overlapping and blurred; part of the stop sound is due to the vowel transition. Liberman et al (1968) saw this difference between the auditory and visual systems as so fundamental that they asserted "no amount of training will cause an appropriate speech decoder to develop for a visual input" (p. 131). Victor Zue has proven their appraisal wrong, but he has also shown that their analysis of the source of difficulty may be correct: much of his ability is based on his knowledge of coarticulation (context-dependent) effects.

Why should context dependency and its associated segmentation problem produce learning difficulties? According to Liberman et al (1968), the nature of the speech code is such that while the auditory system has developed to deal with its temporal properties, the visual system is not capable of processing it in a spatial layout. Yet Victor Zue's performance demonstrates that it can be accomplished. The question then is why, from a perceptual learning point of view, context-dependent features are difficult to identify. One suggestion comes from recent work by Biederman and Shiffrar (1987) on chick-sexing. Biederman and Shiffrar (1987) demonstrated that for the skill of determining the gender of day-old chicks, training time could be drastically reduced by identifying non-accidental distinguishing features. Chick-sexing reportedly takes several years of essentially trial and error practice to achieve high proficiency. By identifying simple invariant features, Biederman and Shiffrar were able to reduce these years of training to a simple rule for finding a distinguishing contour. Although they didn't show why learning was originally so difficult, Biederman and Shiffrar hypothesized that the critical distinguishing features were obscured by their small size and by being embedded in other parts. In such cases, they concluded, it is better to provide instruction which points out the features than to hope they will be discovered by the learner.

The same causes of difficulty may apply to the reading of speech spectrograms. The context-dependent nature of the speech signal causes the visual system to break up the display in inappropriate places. Additionally, cognitive processes may be more likely to group certain parts together into objects and restrict attention to these object units (Ceraso, 1985; Kahneman, 1973). This may produce search difficulties if features required to identify one pattern are spread across different objects. An otherwise noticeable distinguishing feature may be difficult to discover because it is in another "part." This hypothesis is examined in the experiments which follow.

The question of interest to the present work is whether the difficulty of learning spectrogram reading is produced by context-dependent relations among visual features. To enable experimental manipulation of the relations of interest, pseudo-spectrograms were

used. A computer program generated these pseudo-spectrograms based on feature descriptions and interaction rules of real speech spectrograms (Zue, unpublished). A general resemblance to actual spectrograms was maintained.

The patterns used in the experiments were composed of two- or three-phoneme syllables in a vowel-consonant or consonant-vowel-consonant order. Examples of the pseudo-spectrograms used in Experiment 1 are shown in Figure 1. The consonants used were the stop consonants /b/, /p/, /t/, /k/, /d/, and /g/. The vowels used were /i/ as in "beet," /u/ as in "boot," /ae/ as in "bat," /e/ as in "bait," /ɔ/ as in "bought," and /o/ as in "boat." Vowel patterns were quite similar to each other and appeared as wide striated areas with two dark formants (F1 and F2) and one lighter formant (F3). Vowels differed from each other by their width and the height of their three formants.

The purpose of the two experiments described below is, first, to demonstrate that a context-dependent discrimination (i.e., one whose features cross an object boundary) can produce learning difficulty in a pseudo-spectrogram reading task; and second, to look at what contribution segmentation, as distinguished from other factors such as salience, makes to that difficulty.

Experiment 1

To examine the difficulty of learning a context-dependent discrimination, a task was set up to compare the learning of three pairs of consonants. These pairs were /b/-/p/, /t/-/k/, and /d/-/g/. Because the objective was to look at within-pair discriminations, between-pair discriminations were made simple by giving members of the same pair similar widths, but members of different pairs very different widths. Hence, /b/ and /p/ were both very thin, /t/ and /k/ were both wide and /d/ and /g/ were both of medium width. Within-pair discriminations were of three types: multiple cue, single cue, and single context-dependent cue. The consonants /b/ and /p/ differed from each other in texture, shape, and width, and could be distinguished on any of these dimensions. The consonants /t/ and

/k/ could be reliably distinguished only by a single cue. They had the same shape, width, and texture, but a different number of formants. The consonants /d/ and /g/ could also be distinguished only by a single cue, but this cue could not be found by looking at the consonant pattern itself. The shape, width, and texture of /d/ and /g/ were identical and the only way to tell them apart was by their influence on an adjacent vowel. All of the consonants, except /g/, made the second and third formants of an adjacent vowel curve slightly downward at the consonant-vowel boundary. The consonant /g/ made the second and third formants curve toward each other and meet at the consonant-vowel boundary (velar pinch).

The prediction for the experiment was that the context-dependent discrimination would be more difficult to learn than either the single or multiple cue discriminations.

Method

Subjects

Ten subjects were recruited from the University of Pittsburgh. The subjects received credit towards an introductory psychology class and \$10 for their participation.

Apparatus

The pseudo-spectrogram patterns were shown to the subjects on the high resolution display screen of a XEROX 1108 computer. Subjects responded by using a mouse to make selections from a screen menu. The computer collected the subjects' responses and provided accuracy feedback to them.

Materials

The pseudo-spectrogram patterns were generated by a computer program as screen bitmaps. The patterns were 346 X 346 pixels and measured 10 cm X 10 cm on the display screen. The phoneme patterns were drawn from descriptions which mapped a random

texture of a particular shade of grey to different regions of the space the pattern was to occupy. The patterns were drawn as lines of these small texture patterns, the length of which was predetermined except when a line bordered a blank area. In that case the ending point of the line was set to a random number within 10 pixels (3 mm) of its predetermined ending point. Texture and line-length randomization thus provided a small amount of random variability in reappearances of the same phoneme.

The patterns for the phonemes /b/ and /p/ were thin long lines of either a more striated (/b/) or more random (/p/) texture. For the phonemes /t/ and /k/, the patterns were a background of random texture with either a single dark area (for /k/) or two dark areas (for /t/). Because the descriptions for the background textures of /t/ and /k/ were identical, the only reliable way of distinguishing between them was the presence of the extra dark area in /t/. The phonemes /d/ and /g/ appeared as long striated patterns before a vowel and as short striated patterns with two appendages after a vowel, but because their descriptions were identical, the only reliable way to distinguish between them was by the convergence or lack of convergence of the formants in the adjacent vowel. Vowel patterns appeared as a striated uniform background with two dark lower bars below a lighter bar. Vowels could be discriminated by the amount of space between their formants. When vowel formants were curved by the presence of an adjacent /g/, only the center of the pattern could be used to determine the real distance between formants.

Design

Subjects participated in four one-hour sessions held on consecutive days except for one of the subjects who participated in only three sessions but learned all of the discriminations. The spectrogram patterns the subjects saw were all possible consonant-vowel-consonant combinations of the consonants /b/, /p/, /t/, /d/, /g/, /k/, and the vowels /i/, /e/, /æ/, /ɔ/, /o/, /u/. The total number of different combinations was 216. Half of these "words" (108) were used in each session so that after four sessions the subjects saw each word pattern only twice. To control for the frequency of seeing

each phoneme, the words were blocked into groups of six in which each consonant appeared once in prevocalic and postvocalic form, and each vowel appeared once. A subject saw 18 such blocks in each session. Before each session, the order of the words within each block and the order of the blocks within the session were randomized.

Procedure

Subjects were tested individually. A subject was seated in front of the computer and shown how to use a mouse to choose a letter response from a screen menu. The experimenter then briefly explained about spectrograms and told the subject that his or her task was to learn which letters were represented by each pattern. It was made clear, however, that the task was a visual one, and the subjects were discouraged from using strategies based on the sound properties of the phonemes, such as stress or pitch.

When the experiment began, a pseudo-spectrogram pattern appeared in the center of the display screen and remained there until a response was given. Immediately after the pattern's appearance, the message "Think about your answer..." appeared above the pattern in a message box. Because of program differences, three of the subjects saw this message on the screen for 20 seconds, while for the remaining subjects the message remained on the screen for only 3 seconds. This difference was not expected to influence the results because most responses, especially early in the experiment, required more than 20 seconds. Next, a menu appeared on the screen along with the message "Click on the first sound in the word." The menu contained a list of the consonant responses and an example word in which the consonant is used. After a subject selected one of the consonants, a vowel menu appeared with the message "Click on the second sound in the word." Once the vowel was selected, the consonant menu reappeared for the third response. After the subject made the final response, the program provided feedback. If all three responses were correct, the message "That's correct" was displayed in the message box. Otherwise, the message "That's wrong" was displayed along with the correct answer. The pseudo-spectrogram pattern remained on the screen for five seconds after feedback was given. The

subjects were allowed to take a short break halfway through the session.

Shortly after the beginning and toward the end of each session, the experimenter turned on a tape recorder and asked the subject to continue with the next six trials but describe verbally what he was looking at in the pattern and how he decided what to respond.

Results and Discussion

A subject was considered to have learned a consonant pair if he or she responded correctly to four consecutive trial blocks (8 problems) with one allowed error on the third or fourth block. The learning point was taken as the first of the four blocks. Not all of the subjects were able to learn all three consonant discriminations within the allotted time. Of the 10 subjects, 9 learned the /b/-/p/ distinction, 6 learned the /t/-/k/ distinction, and 2 learned the /d/-/g/ distinction. McNemar's exact test for correlated proportions indicated that significantly more subjects learned the /b/-/p/ distinction than the /d/-/g/ distinction ($p < .02$), but the test of whether more people learned the /t/-/k/ distinction than learned the /d/-/g/ distinction was not significant ($p = .10$).

A matched pairs sign test was used to test whether the learning points for the /b/-/p/ and /t/-/k/ distinctions were earlier than for the /d/-/g/ distinction. Unlearned distinctions were considered to have a learning point of at least 73 (i.e., one greater than the last block). If two distinctions were unlearned, the learning points were considered to be tied. Using this procedure, the /b/-/p/ and /t/-/k/ distinctions were found to have been learned at an earlier point than the /d/-/g/ distinction ($p < .01$ and $p < .02$ respectively).

To obtain a measure of how much earlier the single- and multiple-cue distinctions were learned, it was necessary for the subjects to have learned to distinguish at least two of the three consonant pairs. Four subjects failed to meet this criterion and were not included in the measure. Of the six remaining subjects, only two learned the /d/-/g/ distinction. For the others, the learning point was

estimated as 73. Because this value underestimates the true learning point, the measure of when the /d/-/g/ distinction was learned is conservative. Based on this measure, the mean number of trial blocks required for subjects to learn each consonant pair discrimination is provided in Table 1. According to these estimates, the /d/-/g/ distinction appears to require a considerably greater amount of learning time than either the /b/-/p/ or /t/-/k/ distinctions (approximately 40 additional blocks).

Consonant Distinction

	<u>Multiple Cue</u> <u>/b/-/p/</u>	<u>Single Cue</u> <u>/t/-/k/</u>	<u>Context Cue</u> <u>/d/-/g/</u>
Mean	20.17	29.17	66.17
Standard deviation	17.81	23.26	12.17
Number of estimated points	0	0	4

Table 1: Mean number of trial blocks to reach learning criterion for each consonant distinction.

These results suggest that a context-dependent discrimination can be difficult to learn. Fewer subjects were able to learn the /d/-/g/ discrimination in the allotted time. The test on proportion of learners for each distinction showed that significantly more people learned the multiple-cue distinction than the context-dependent one. The difference between the proportion who learned the single-cue distinction and the context-dependent one, though not significant, was large (.60 vs .20). For those subjects who did learn the context-dependent discrimination (or who were optimistically presumed to be about to learn it when the experiment ended), learning took longer than for either the multiple cue or the single cue discrimination. These findings suggests that having to discover a context-dependent discrimination could account for some of the difficulty encountered in acquiring the skill of speech spectrogram reading.

However, these results must be viewed with caution. The experiment examined learning of a realistic and complex pattern, and likely confounded several factors with the context-dependent vs non-context-dependent comparison. These factors must be ruled out before learning difficulty can be unambiguously assigned to the context-dependent manner in which the stimulus is segmented. One such factor is cue salience. It may simply have been harder for the subjects to notice the formant curving cue than the other cues. This explanation is unlikely given that 8 of the 10 subjects mentioned in their verbal reports that there was something unusual about the appearance of the formants (i.e., that they were curved or straight). Nevertheless, salience differences must be ruled out. Another confounding factor is whether task demands, rather than segmentation difficulty, made the /d/-/g/ distinction difficult to learn. Subjects may have noticed the formant curving cue, but because they also were required to learn the identity of the vowel, may have tried to use formant curving to distinguish among the different vowels. This may have "used up" the cue, making it unavailable for use in distinguishing the consonants. There is support for this possibility in the verbal reports made by several subjects who mentioned the formant curving in conjunction with vowel discriminations. These two possible alternative explanations are examined in Experiment 2.

Experiment 2

In Experiment 2, the goal was to try to determine whether the learning difficulty observed in Experiment 1 was due to context-dependent segmentation, to some other factor such as salience or task demands, or to some interaction of these factors. Segmentation, in this context, refers to how the cognitive system divides a pattern into objects. Segmentation was manipulated by having two cues occur within the same object or by splitting them between two objects. Salience is how noticeable the features are. This was measured by having a separate group of subjects circle the parts in the spectrogram patterns used in this experiment. It was also controlled for in the experimental design by having different groups of subjects learn each distinguishing cue both as a between-object cue and as a within-object

cue. Finally, task demands refer to whether the subject was to treat the different phonemes as separate parts in making a response. In this experiment subjects made only a single response to the whole pattern, but an attempt to vary task demands was made through instructional bias.

Method

Materials

The pseudo-spectrogram patterns used in Experiment 2 were similar to those used in Experiment 1, but to control for all of the independent variables, several changes were made. First, the patterns consisted of only two phonemes: a vowel-like pattern, followed by a consonant-like pattern. The vowel patterns were either thin (T) or wide (W), and had formants which were either straight (S) or curved (C) and either high (H) or low (L) in frequency (/i/ vs /ae/). Consonant patterns could be large (L) or small (S) and had either one (O) or two (T) formants. Formants appeared as dark spots on the large consonants and as protrusions on the small consonants. Figure 2 shows some examples of these patterns. The pseudo-spectrogram patterns were generated in the same way as those in Experiment 1; the 32 vowel-consonant combinations were drawn 8 times for a total of 256 patterns.

To assess the salience of the patterns' visual features, a group of 15 subjects (not the same as those in the learning task) were given a stack of the 32 different patterns and asked to circle the "important parts." The results of this circling task are given in Table 2. Of relevance to the present experiment is the finding that the subjects circled the vowel formants an average of 98% of the time, while circling the consonant formants an average of only 76% of the time. Furthermore, the subjects tended to circle curved vowel formants as a single part (67% of the time), and straight vowel formants as separate parts (83% of the time). The first consonant formant was circled more often than the second (81% vs 68%), and formants in the large consonants were circled more often than formants in the small consonants (90% vs 59%). Hence, some of the difference in salience

between vowel formants and consonant formants may be due to difficulty seeing the small consonant formants as distinct parts.

<u>Feature</u>	<u>Proportion</u>
Whole Vowel	.13
1st Vowel Formant	.97
2nd Vowel Formant	.99
3rd Vowel Formant	.98
All Other Vowel Features	.22
Whole Consonant	.33
1st Consonant Formant	.83
2nd Consonant Formant	.69
All Other Consonant Features	.29

Table 2: Proportion of times a feature was circled in part-circling task.

Design

The goal of the experiment was to assess whether a within-object cue would be learned more readily than a between-object cue. To avoid confounding the type of cue (formant curving or number of formants) with the location of the cue (within or between objects), each cue type was learned as both a within-object cue and as a between-object cue. Because this could not be manipulated within subjects, an incomplete blocks design was used. Each subject provided two observations from the 2 X 2 (Cue Type X Cue Location) design, and a block of two subjects with complementary conditions constituted a single replication of the design. This confounds the Cue Type X Cue Location interaction with subjects, but by running enough replications, this effect could be analyzed as a between block factor.

One additional factor, instruction, was also included as a between block factor. One half of the blocks received neutral instructions which asked them to learn to associate the whole pattern

with a response, the other half received biasing instructions which asked them to learn the half of the pattern containing the within-object cue. The within and between block designs made up four conditions: Neutral Instructions, Curve-Within (NCW); Neutral Instructions, Curve-Between (NCB); Biased Instructions, Curve-Within (BCW); and Biased Instructions, Curve-Between (BCB). The Curve-Within/Curve-Between distinction refers to the type of rules subjects were to learn. Table 3 shows these rules for each condition.

Cons.	Condition (Instructions- Curve location)	Left Pattern	Right Pattern
/g/ /d/ /k/ /t/	Neutral-Within (NCW)	Curved,Thin Straight,Thin Wide Wide	One formant Two formants
/g/ /d/ /k/ /t/	Neutral-Between (NCB)	Curved Straight	Small Small Large and One formant Large and Two formants
/g/ /d/ /k/ /t/	Biased-Within (BCW)	Curved,Thin Straight,Thin Wide Wide	One formant Two formants
/g/ /d/ /k/ /t/	Biased-Between (BCB)	Curved Straight	Small Small Large and One formant Large and Two formants

Table 3: Rules for discriminating patterns in Experiment 2

The Curve-Within groups learned the formant curving cue as a within-object cue and the number of formants cue as a between-object cue; the Curve-Between groups learned the formant curving cue as a between-object cue and the number of formants cue as a within-object cue.

Subjects participated in a single two hour session. The pseudo-spectrogram patterns the subjects saw were all possible vowel-consonant combinations as described above. To control for the frequency of seeing each phoneme, the patterns were grouped into blocks of eight in which each consonant appeared twice and each vowel appeared once. Before each session, the order of the patterns within each block, and the order of the blocks within the session were randomized for each subject.

Procedure

Subjects were tested individually. Each subject was seated in front of the computer and shown how to use a mouse to choose a letter response from a screen menu. Then the instructions for the experiment were displayed on the screen. Subjects in the neutral conditions were told their task was to learn to identify which pattern was displayed; subjects in the biased conditions were told to identify the left (or right) pattern. To ensure that the subjects in the biased condition read the instructions, they were asked to identify which half (left or right) of the pattern they were to learn. If they were incorrect, the instructions reappeared on the screen.

The experiment began with a pseudo-spectrogram pattern appearing in the center of the display screen. The message "Think about your answer..." appeared in a message box above the pattern for 3 seconds. Then a menu appeared on the screen along with the message "Click on the first sound in the word." The menu contained a list of four responses: /t/, /k/, /d/, and /g/. After the subject made a response, the program provided feedback. If the response was correct, the message "That's correct" was displayed in the message box. Otherwise, the message "That's wrong" was displayed along with the correct answer. Once feedback was given, the pseudo-spectrogram

pattern remained on the screen for 10 seconds before being replaced by the pattern for the next trial. Every 32 trials, the subject was allowed to take a short break before continuing.

After the session, the experimenter turned on a tape recorder and asked the subject to identify 8 patterns and describe what she looked at in the pattern and how she decided what to respond.

Subjects

Forty-eight introductory psychology students from the University of Pittsburgh participated for course credit. Two subjects, both from the Neutral-Curve-Between condition, were replaced: one quit the session early, the other hadn't slept for 48 hours prior to the experiment session and showed no learning. The remaining subjects were randomly assigned to the four conditions with the constraint of obtaining 9 full or partial learners (as described below) in each condition.

Results and Discussion

As in Experiment 1, subjects had considerable difficulty learning both the within and between object distinctions. A subject was considered to have learned a distinction when correct responses were made on two consecutive blocks (8 problems) with one allowed error on the second block (two subjects were also considered to have learned a distinction on their final block if the final block was correct and they gave the correct rule for the distinction in their post-session interview). By this criterion, the 48 subjects fall into three categories: full learners, non-learners, and partial learners. Full learners were those who learned both the between and within object distinctions; non-learners learned neither distinction; partial learners were those who only learned one of the two distinctions. Table 4 summarizes how the subjects performed. Eighteen subjects were full learners, twelve were non-learners, and eighteen were partial learners. Of the partial learners, 13 learned only the within rule and 5 learned only the between rule. Of the non-learners, one was from the NCB condition, two from the BCW condition, and nine from the NCW condition.

Discriminations learned	NCW	NCB	BCW	BCB
Both discriminations	4	9	1	4
One discrimination				
Within rule only	5	0	7	1
Between rule only	0	0	1	4
Neither discrimination	9	1	2	0

Table 4: Discriminations learned, by condition

A matched pairs sign test was used to test the main effects of Cue Location and Cue Type for those subjects who were full or partial learners. For partial learners, the learning point of the unlearned distinction was considered to be at least 17 (the last trial block plus one). By this test, the main effect of Cue Location was not significant ($z=1.39$, $p<.09$), but the main effect of Cue Type was significant ($z=4.18$, $p<.001$). The subjects learned the formant curving cue before the number of formants cue significantly more often than they learned them in the reverse order. To test the interaction of Cue Type X Cue Location, each subject's performance was categorized according to its sign. A chi-square test of independence revealed that the interaction was significant ($\chi^2(2)=19.35$, $p<.001$). Formant curving was learned first as a within-object cue just as often as it was learned first as a between-objects cue, but the number of formants cue was learned first as a within-object cue more often than as a between-objects cue.

To obtain a measure of when the distinctions were learned, the learning point for unlearned distinctions was estimated as the 17th block. This value underestimates the true learning block and makes the measure conservative. Most of these estimations were made for the between-object distinction when it involved the number of formants cue. This is also consistent with the observation that an unusually large number of non-learners were found in the conditions which required learning this distinction (the NCW and BCW conditions). Making these estimations, the mean learning block for each distinction

and condition was calculated. These values are given in Table 5. The measures indicate that the number of formants cue was learned at least five blocks earlier as a within-object cue than as a between-object cue, but the formant curving cue was learned at about the same point for both locations.

<u>Cue Type</u>	<u>Cue Location</u>	
	Within	Between
Number of formants		
Mean	10.28	15.56
Standard deviation	5.04	2.59
Number of estimated points	4	12
Formant curving		
Mean	8.72	7.44
Standard deviation	4.23	3.75
Number of estimated points	1	1

Table 5: Mean number of trial blocks to reach learning criterion for each consonant distinction

These results suggest that lack of salience may play an important part in making this type of skill difficult to learn. The sign test demonstrates that the formant curving cue was more often learned before the number of formants cue, and the estimates of learning points shows that the formant curving cue was learned at least 4 blocks earlier, on average. The cause of this difference is likely to be cue salience. In the part circling task, more subjects circled the vowel formants than the consonant formants, suggesting that the vowel formants are more salient. The effect of salience, however, does not explain the learning difficulty observed in the first experiment. In Experiment 1, number of formants as a within-object cue was learned sooner and more often than the formant curving cue as a between-object cue. If this were due to salience, then we should have found that the number of formants cue was learned sooner than the formant curving cue in Experiment 2.

Nor can segmentation by itself account for the observed learning difficulty. Cue Location was not significant, and even the interaction of Cue Type and Cue Location does not produce a simple explanation. Context-dependent segmentation does appear to produce learning difficulty, but this effect may be restricted to cues of lower salience. The chi-square test on the interaction of Cue Type and Cue Location showed that more subjects learned the formant curving cue before the number of formants cue when the number of formants cue was a between-objects cue, but when the number of formants cue was a within-objects cue, the order of learning was indifferent to cue type. Thus, difficulty due to cue location was found for the less salient number of formants cue but not for the more salient formant curving cue. However, the degree of impairment for less salient cues appears to be substantial. More non-learners (11 vs 1) and within-rule-only learners (12 vs 1) were reported in the conditions which required learning the number of formants cue as a between-objects cue. Additionally, the conservative estimate of learning points indicates that this cue was learned at least five blocks later as a between- than as a within-object cue.

Yet segmentation does not explain the learning difficulty observed in the first experiment. In Experiment 1, the formant curving cue as a between-object cue was found to be much harder to learn than the number of formants cue as a within-object cue. This finding was not replicated in the second experiment. In fact, the opposite was found. Neither salience nor segmentation can account for this difference because neither was changed between the two experiments. The only major change was the learning task.

Presumably, the reason the formant curving cue was difficult to learn in Experiment 1 was the vowel response required in that task. This was not manipulated in the second experiment, so it is impossible to be certain. It is interesting to note, however, that the difficulty disappeared when the vowel identification task was eliminated in Experiment 2. Unfortunately, the manipulation of instructional bias in this experiment was too weak to clarify this question. Half of the subjects were instructed to "learn to identify the right [or "left"] hand

part" of the pattern, but in post-experiment interviews several admitted to ignoring these instructions. Instructional bias did not significantly interact with either Cue Type or Cue Location ($\chi^2(2)=3.31$, $p>.10$, $\chi^2(2)=3.82$, $p>.10$, respectively). Future research should determine whether task demands cause the difficulty observed in Experiment 1 by more strongly manipulating task demands within a single experiment.

General Discussion

The two experiments presented here point to several factors which can affect the difficulty of learning to read speech spectrograms. The original hypothesis, that learning difficulty was caused by context-dependent relations created by the way the visual system segments spectrogram patterns, has been shown to be too simple. Learning difficulty for this skill may be affected by the interaction of segmentation with cue salience and task demands. Segmentation was shown to have a considerable influence on difficulty, but this influence may be restricted to less salient cues. Segmentation may also be influenced by the demands of the learning task. Although the experiments did not demonstrate this, it is likely that the type of response required by the learning task influences task difficulty. The following discussion examines in more detail why segmentation might interact with these factors.

The interaction of segmentation with cue salience can be explained by assuming that whatever learning difficulties are produced by segmentation can be overcome by a highly salient cue. Salience has long been known to influence hypothesis selection in discrimination learning tasks (Trabasso & Bower, 1968). Highly salient cues are likely to be tried first as hypotheses. If the effect of segmentation is to make certain cues less available for selection as hypotheses, then it is easy to understand why a high degree of salience would overcome this effect. This explanation is supported by the results of the second experiment reported here, in which the mean learning block was about the same for all distinctions except for the condition when the less salient number of formants cue was a between-objects cue. When the formant curving cue was a between-

objects cue, its highly salient nature made it available for attention anyway.

Although neither experiment directly manipulated task demands, the difference between the results of the two experiments suggests that the type of response the subjects were required to give was also important. In the first experiment, where the subjects were required to respond to both consonants and vowels, they had difficulty learning the highly salient formant curving cue as a between object cue. In the second experiment, where subjects made only a single response to the whole pattern, formant curving was no more difficult to learn as a between-object cue than as a within-object cue. Since subjects in Experiment 1 reported using formant curving to distinguish the vowel responses, it seems likely that including the vowel response made it more difficult to notice the relevance of the formant curving to the consonant distinction, perhaps in the following way. A subject might select the cue as a hypothesis for vowel identification. When this hypothesis was disconfirmed, the hypothesis may have become less likely to be selected immediately again. If the formant curving cue was selected as relevant for vowel discrimination because of the way that spectrograms are segmented visually, it might be less available for part of a consonant discrimination. In the second experiment, when the vowel identification task was eliminated, subjects were more able to learn formant curving as a between object cue.

Task demands may also have increased learning difficulty by reinforcing any existing segmentation biases. If subjects were required to make two responses to a pattern, they may have been more likely to see the pattern as two distinct parts, and possibly to assign one response to one part, and the other response to the remaining part. This may have enhanced any existing bias against crossing part boundaries. This hypothesis can be tested only by future research.

The main conclusion of the present research is to confirm the influence of segmentation on learning difficulty in speech spectrogram reading. Although segmentation was not found to be the sole determiner of such difficulty, in combination with other stimulus and task variables it appeared to have a substantial influence. One way of

thinking about the effect of segmentation is as a within-object search bias. People may be biased toward searching within an object's part boundaries (contour) for discriminating features, before considering features outside those boundaries. This bias, however, can be overridden by a highly salient feature in another part. The learning task is also important to the within-object search bias. If a feature can be used as a within-object cue, then it may be less likely to be considered as a between-object cue. Such factors may have led the subjects in Experiment 1 to believe incorrectly that formant curving indicated vowel identity, and may have impaired their ability to associate it with consonant identity.

The existence of a within-object search bias is consistent with several theories of visual attention. According to the view taken by Kahneman (1973; Kahneman & Henik, 1981) and Ceraso (1985), attention to a visual scene is allocated by object units. According to Kahneman's (1973) model of attention and perception, preattentive visual processes divide a display into units according to stimulus properties and simple grouping rules (such as Gestalt rules). These units are given figural emphasis (attention) based on factors such as figure-ground relations, features which make something STAND OUT, and intention. Units which receive this attention are then matched against memory structures to test for recognition. Visual search involves the intentional switching of figural emphasis from object to object, or the attraction of figural emphasis based on a feature (either stimulus or response selected) which distinguishes the target. According to the results of the experiments presented above, the features of a target phoneme unit are more likely to be considered than features of other phonemes, unless those other features are highly salient. This result may be due to the way attention is allocated to a whole part unit. If whole phonemes are attended as wholes, then the features within the attended phoneme will receive figural emphasis and be further processed as potential hypotheses. However, if a highly salient feature, one which draws attention to itself, is in a neighboring phoneme, it may be included in processing and may even be selected earlier as a hypothesis. According to this attention-by-parts view, the within-object search bias may be the result of normal attention allocation policy within the visual system.

A within-object search bias is also consistent with recent suggestions that preferences and heuristics are required to restrict the amount of search involved in concept learning (Michalski, 1983; Medin, Wattenmaker & Michalski, 1987). This view is not inconsistent with the attention-by-parts hypothesis, but emphasizes the functional role of such a bias in the learning process. In complex visual environments, ordered search for important features (even salience ordered search) is too resource consuming to be viable. Rather, preferences for certain features or locations are required to restrict the scope of search. Restricting the search for a discriminating feature to the area within the object boundaries of a part is a sensible heuristic. In our normal visual perception, objects are classified or discriminated by features within their own object boundaries. Only in certain artificial environments, such as speech spectrograms or x-ray pictures, are context-dependent relations set up by visual segmentation. In such environments, what is normally a useful heuristic actually hinders search rather than aiding it.

In the second experiment, what was observed was not a facilitating effect for a within-object cue, but an increased difficulty for locating a between-object cue. Cues with low salience can be fairly easily located when they are within the same object, but when a low-salience cue must be found in a nearby object, learning difficulty is increased, probably by a tendency to retry discarded within-object hypotheses. This result has important implications for speech spectrogram reading. First, it explains at least part of the enormous difficulty in learning the skill of speech spectrogram reading. In spectrogram reading, the large variability in the appearance of phonemes means that the salience of most features is likely to be quite low. Also, it is important to learn spectrogram patterns at the individual phoneme level. Hence, the narrow focus induced by the task should be expected to increase the within-object search bias and impair discovery of context-dependent features.

Some individuals, too, might be more affected by a search bias than others. For some, it may only slow down search, with the low-salience context-dependent feature found only after within-object


features have been searched. For others, it may mean the complete abandonment of search after a within-object search has failed. Such differences depend on an individual's repertoire of strategies and learning history. Fortunately for students of spectrogram reading, Victor Zue has identified many of these features, so they do not have to be discovered anew.

In most visual environments and for most perceptual skills, a within-object bias is helpful. It restricts the amount of search required for learning. However, for other environments and skills, such as speech spectrogram reading, radiology, and passive sonar reading, where visual objects and real objects do not directly correspond (Lesgold et al, 1988; Liberman et al, 1968; Smith, 1982), it becomes a source of learning difficulty. Overcoming such search biases may be an important part of learning for these skills.


References

- Biederman, I., & Shiffrar, M. M. (1987). Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 640-645.
- Ceraso, J. (1985). Unit formation in perception and memory. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 19, (pp. 179-210). New York: Academic Press.
- Chase, W.G. & Simon, H.A. (1973). The mind's eye in chess. In W.G. Chase (Ed.), *Visual information processing* (pp. 215-281). New York: Academic Press.
- Cole, R. A., Rudnicky, A. I., Zue, V. W., & Reddy, D. R. (1980). Speech as patterns on paper. In R. A. Cole (Ed.) *Perception and production of fluent speech* (pp. 4-50). Hillsdale, NJ: Erlbaum.
- Cole, R. A. & Zue, V. W. (1980). Speech as eyes see it. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 475-494). Hillsdale, NJ: Erlbaum.
- Greene, B.G., Pisoni, D.B., & Carrell, T.D. (1984). Recognition of speech spectrograms. *Journal of the Acoustical Society of America*, 76, 32-43.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kahneman, D., & Henik, A. (1981). Perceptual organization and attention. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 181-211). Hillsdale, NJ: Erlbaum.
- Lesgold, A., Robinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing X-ray pictures. In M. T. H. Chi, R. Glaser, & M. Farr (Eds.), *The nature of expertise*. Hillsdale, NJ: Erlbaum.
- Levine, M. (1975). Hypothesis behavior by humans during discrimination learning. In M. Levine (Ed.), *A cognitive theory of learning: Research on hypothesis testing* (pp. 181-190). Hillsdale, NJ: Erlbaum. (Adapted from *Journal of Experimental Psychology*, 1966, 71, 331-338).


- Liberman, A. M., Cooper, F. S., Shankweiler, D. P., Studdert-Kennedy, M. (1968). Why are speech spectrograms hard to read? *American Annals of the Deaf*, 113, 127-133.
- Medin, D. L., Wattenmaker, W. D., & Michalski, R. S. (1987). Constraints and preferences in inductive learning: An experimental study of human and machine performance. *Cognitive Science*, 11, 299-339.
- Michalski, R. S. (1983). A theory and methodology of inductive learning. *Artificial Intelligence*, 20, 111-161.
- Potter, R.K., Kopp, G.A., & Green, H.C. (1947). *Visible speech*. New York: D. Van Nostrand Co. Inc.
- Smith, M. J. (1982). Human factors research in passive sonar operating. *Proceedings of the international conference on man machine systems*. London: Institute of Electrical Engineers.
- Squire, L.F. (1988). *Fundamentals of radiology*. Cambridge, MA: Harvard University Press.
- Trabasso, T. & Bower, G. H. (1968). *Attention in learning*. New York: Wiley.
- Triesman, A. (1986). Properties, parts, and objects. In K.R. Boff, L. Kaufman, and J.P. Thomas (Eds.), *Handbook of perception and human performance*. Volume II: *Cognitive processes and performance*. New York: Wiley.
- Zue, V. W. (unpublished). *Notes on spectrogram reading*. Cambridge, MA: Massachusetts Institute of Technology, Department of Electrical Engineering & Computer Science and Research Laboratory of Electronics.



[b i p]




[d æ g]




[t ø k]


Figure 1. Examples of pseudo-spectrograms used in Experiment 1.




Thin-Straight-High Large-One



Wide-Straight-Low Large-Two



Thin-Curved-High Small-One



Wide-Curved-Low Small-Two

Figure 2. Examples of pseudo-spectrograms used in Experiment 2.

University of Pittsburgh/Lesgold

Dr. Edith Ackermann
Media Laboratory
E15-311
20 Ames Street
Cambridge, MA 02139

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Stephen J. Andriole, Chairman
Department of Information Systems
and Systems Engineering
George Mason University
4400 University Drive
Fairfax, VA 22030

Prof. John Annett
University of Warwick
Department of Psychology
Coventry CV4 7AL
ENGLAND

Dr. Gary Aston-Jones
Department of Biology
New York University
Rm. 1009 Main Bldg.
100 Washington Square East
New York, NY 10003

Dr. Patricia Baggett
School of Education
610 E. University, Rm 1302D
University of Michigan
Ann Arbor, MI 48109-1259

Dr. Gautam Biswas
Department of Computer Science
Box 1688, Station B
Vanderbilt University
Nashville, TN 37235

Dr. Lyle E. Bourne, Jr.
Department of Psychology
Box 345
University of Colorado
Boulder, CO 80309

Dr. Robert Calfee
School of Education
Stanford University
Stanford, CA 94305

Dr. Joseph C. Campione
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Dr. Jaime G. Carbonell
Computer Science Department
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. John M. Carroll
IBM Watson Research Center
User Interface Institute
P.O. Box 704
Yorktown Heights, NY 10598

Dr. David E. Clement
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Jere Confrey
Cornell University
Dept. of Education
Room 490 Roberts
Ithaca, NY 14853

Dr. Lynn A. Cooper
Department of Psychology
Columbia University
New York, NY 10027

Dr. Thomas E. DeZern
Project Engineer, AI
General Dynamics
PO Box 748/Mail Zone 2640
Fort Worth, TX 76101

Dr. Ronna Dillon
Department of Guidance and
Educational Psychology
Southern Illinois University
Carbondale, IL 62901

University of Pittsburgh/Lesgold

Dr. J. Stuart Donn
Faculty of Education
University of British Columbia
2125 Main Mall
Vancouver, BC CANADA V6T 1Z5

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
(12 Copies)

Dr. John Ellis
Navy Personnel R&D Center
Code 51
San Diego, CA 92252

ERIC Facility-Acquisitions
2440 Research Blvd, Suite 550
Rockville, MD 20850-3238

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Campus Box 345
Boulder, CO 80309-0345

Dr. Marshall J. Farr, Consultant
Cognitive & Instructional Sciences
2520 North Vernon Street
Arlington, VA 22207

CAPT J. Finelli
Commandant (G-PTE)
U.S. Coast Guard
2100 Second St., S.W.
Washington, DC 20593

Dr. Norman Frederiksen
Educational Testing Service
(05-R)
Princeton, NJ 08541

Dr. Alinda Friedman
Department of Psychology
University of Alberta
Edmonton, Alberta
CANADA T6G 2E9

Dr. Robert M. Gagne
1456 Mitchell Avenue
Tallahassee, FL 32303

Psychology Department
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Marvin D. Glock
101 Homestead Terrace
Ithaca, NY 14856

Dr. Dwight J. Goehring
ARI Field Unit
P.O. Box 5787
Presidio of Monterey, CA 93944-5011

Dr. Sherrie Gott
AFHRL/MOMJ
Brooks AFB, TX 78235-5601

Dr. Wayne Gray
Artificial Intelligence Laboratory
NYNEX
500 Westchester Avenue
White Plains, NY 10604

H. William Greenup
Dep Asst C/S, Instructional
Management (E03A)
Education Center, MCCDC
Quantico, VA 22134-5050

Dr. Gerhard Grossing
Atominstitut
Schuttelstrasse 115
Vienna
AUSTRIA A-1020

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08542

University of Pittsburgh/Lesgold

Dr. James Hiebert
Department of Educational
Development
University of Delaware
Newark, DE 19716

Dr. Steven A. Hillyard
Department of Neurosciences, M-008
University of California, San Diego
La Jolla, CA 92093

Ms. Julia S. Hough
Cambridge University Press
40 West 20th Street
New York, NY 10011

Dr. William Howell
Chief Scientist
AFHRL/CA
Brooks AFB, TX 78235-5601

Z. Jacobson
Bureau of Management Consulting
701-365 Laurier Ave., W.
Ottawa, Ontario K1H 5W3
CANADA

Dr. Daniel B. Jones
U.S. Nuclear Regulatory
Commission
NRR/ILRB
Washington, DC 20555

Mr. Paul L. Jones
Research Division
Chief of Naval Technical Training
Building East-1
Naval Air Station Memphis
Millington, TN 38054-5056

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
Department of Psychology
University of California
Berkeley, CA 94720

Dr. Ruth Kanfer
University of Minnesota
Department of Psychology
Elliott Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Michael Kaplan
Office of Basic Research
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Milton S. Katz
European Science Coordination
Office
U.S. Army Research Institute
Box 65
FPO New York 09510-1500

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Frank Keil
Department of Psychology
228 Uris Hall
Cornell University
Ithaca, NY 14850

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr.
P.O. Box 704
Yorktown Heights, NY 10598

Dr. Thomas Killian
AFHRL/OT
Williams AFB, AZ 85240-6457

Dr. J. Peter Kincaid
Army Research Institute
Orlando Field Unit
c/o PM TRADE-E
Orlando, FL 32813

Mr. David A. Kobus
Naval Health Research Center
P.O. Box 85122
San Diego, CA 92138

University of Pittsburgh/Lesgold

Dr. Sylvan Kornblum
University of Michigan
Mental Health Research Institute
205 Washtenaw Place
Ann Arbor, MI 48109

Dr. Gary Kress
628 Spazier Avenue
Pacific Grove, CA 93950

Dr. Lois-Ann Kuntz
3010 S.W. 23rd Terrace
Apt. No. 105
Gainesville, FL 32608

Dr. Pat Langley
NASA Ames Research Ctr.
Moffett Field, CA 94035

Dr. Robert W. Lawler
Matthews 118
Purdue University
West Lafayette, IN 47907

Dr. John Levine
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Matt Lewis
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Robert Lloyd
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Dr. Don Lyon
P. O. Box 44
Higley, AZ 85236

Dr. William L. Maloy
Code 04
NETPMSA
Pensacola, FL 32509-5000

Dr. Elizabeth Martin
AFHRL/OTE
Williams AFB
AZ 85240

Dr. James G. May
Department of Psychology
University of New Orleans
Lakefront
New Orleans, LA 70148

Dr. Joseph C. McLachlan
Code 52
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. D. Michie
The Turing Institute
George House
36 North Hanover Street
Glasgow G1 2AD
UNITED KINGDOM

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
250 N. Harbor Dr., Suite 309
Redondo Beach, CA 90277

Prof. David Navon
Department of Psychology
University of Haifa
Haifa 31999
ISRAEL

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Librarian
Naval Center for Applied Research
in Artificial Intelligence
Naval Research Laboratory
Code 5510
Washington, DC 20375-5000

Dr. Paul O'Rourke
Information & Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
Pittsburgh, PA 15260

University of Pittsburgh/Lesgold

Dr. James B. Olsen
WICAT Systems
1875 South State Street
Orem, UT 84058

Office of Naval Research,
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Dr. Judith Orasanu
Basic Research Office
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Jesse Orlansky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

Dr. Okchoon Park
Army Research Institute
PERI-2
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Roy Pea
Institute for Research
on Learning
2550 Hanover Street
Palo Alto, CA 94304

Dr. C. Perrino, Chair.
Dept. of Psychology
Morgan State University
Cold Spring La.-Hillen Rd.
Baltimore, MD 21239

Dept. of Administrative Sciences
Code 54
Naval Postgraduate School
Monterey, CA 93943-5026

Dr. Martha Polson
Department of Psychology
University of Colorado
Boulder, CO 80309-0345

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309-0345

Dr. Michael I. Posner
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Robert B. Post
Psychology Department
University of California
Davis, CA 95616

Dr. Daniel Reisberg
Reed College
Department of Psychology
Portland, OR 97202

Dr. Ernst Z. Rothkopf
AT&T Bell Laboratories
Room 2D-456
600 Mountain Avenue
Murray Hill, NJ 07974

Dr. Arthur G. Samuel
Yale University
Department of Psychology
Box 11A, Yale Station
New Haven, CT 06520

Dr. Walter Schneider
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Hans-Willi Schroiff
Psychological Research
Henkel Corp./WSC
Henkelstr. 67
D-4000 Duesseldorf 1 FRG
WEST GERMANY

Dr. Randall Shumaker
Naval Research Laboratory
Code 5510
4555 Overlook Avenue, S.W.
Washington, DC 20375-5000

University of Pittsburgh/Lesgold

Dr. Zita M. Simutis
Chief, Technologies for Skill
Acquisition and Retention
ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Saul Sternberg
University of Pennsylvania
Department of Psychology
3815 Walnut Street
Philadelphia, PA 19104-6196

Dr. Patrick Suppes
Stanford University
Institute for Mathematical
Studies in the Social Sciences
Stanford, CA 94305-4115

Dr. John Tangney
AFOSR/NL, Bldg. 410
Bolling AFB, DC 20332-6448

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1205 Coleman Avenue, Box 580
Santa Clara, CA 95052

Dr. Paul T. Twohig
Army Research Institute
ATTN: PERI-RL
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Zita E. Tyer
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Shih-sung Wen
Department of Psychology
Jackson State University
1400 J. R. Lynch Street
Jackson, MS 39217

Mr. Paul T. Wohig
Army Research Institute
5001 Eisenhower Ave.
ATTN: PERI-RL
Alexandria, VA 22333-5600

Dr. Joseph L. Young
National Science Foundation
Room 320
1800 G Street, N.W.
Washington, DC 20550